The Stata Journal (2016)
16, Number 3, pp. 778–804

Estimation of panel vector autoregression in Stata

Michael R. M. Abrigo
Department of Economics
University of Hawaii at Manoa
Honolulu, HI
and Philippine Institute for Development Studies
Manila, Philippines
mabrigo@mail.pids.gov.ph

Inessa Love
Department of Economics
University of Hawaii at Manoa
Honolulu, HI

Abstract. Panel vector autoregression (VAR) models have been increasingly used in applied research. While programs specifically designed to fit time-series VAR models are often included as standard features in most statistical packages, panel VAR model estimation and inference are often implemented with general-use routines that require some programming dexterity. In this article, we briefly discuss model selection, estimation, and inference of homogeneous panel VAR models in a generalized method of moments framework, and we present a set of programs to conveniently execute them. We illustrate the `pvar` package of programs by using standard Stata datasets.

Keywords: st0455, pvar, pvarfevd, pvargranger, pvarirf, pvarsoc, pvarstable, panel, vector autoregression, VAR, dynamic panel

1 Introduction

Time-series vector autoregression (VAR) models originated in the macroeconometrics literature as an alternative to multivariate simultaneous equation models (Sims 1980). All variables in a VAR system are typically treated as endogenous, although identifying restrictions based on theoretical models or on statistical procedures may be imposed to disentangle the impact of exogenous shocks onto the system. With the introduction of VAR in panel-data settings (Holtz-Eakin, Newey, and Rosen 1988), panel VAR models have been used in multiple applications across fields.

In this article, we briefly review panel VAR model selection, estimation, and inference in a generalized method of moments (GMM) framework and provide a package of programs, which we illustrate using two standard Stata datasets. An early article that examined panel VAR was Love and Zicchino (2006), who made the programs available informally to other researchers.[1] This article introduces an updated package of programs

1. As of February 2016, Love and Zicchino (2006) have been cited in 601 research articles; most use the early version of the package of programs to fit panel VAR models. For example, these programs have been used in studies recently published in the *American Economic Review* (for example, Head, Lloyd-Ellis, and Sun [2014]), *Applied Economics* (for example, Mora and Logan [2012]), the *Journal of Macroeconomics* (for example, Carpenter and Demiralp [2012]), and the *Journal of Economic History* (for example, Neumann, Fishback, and Kantor [2010]), among others.

 st0455

with additional functionality, including estimation by Stata's built-in `gmm` command, which allows for use of all available `gmm` options, addition of exogenous variables to the VAR system, and subroutines to implement Granger (1969) causality tests and optimal moment and model selection criteria (MMSC) following Andrews and Lu (2001), among others.

2 Panel VAR

We consider a k-variate homogeneous panel VAR of order p with panel-specific fixed effects represented by the following system of linear equations,

$$\begin{aligned} \mathbf{Y}_{it} = \mathbf{Y}_{it-1}\mathbf{A}_1 + \mathbf{Y}_{it-2}\mathbf{A}_2 + \cdots\ + \mathbf{Y}_{it-p+1}\mathbf{A}_{p-1} + \mathbf{Y}_{it-p}\mathbf{A}_p + \mathbf{X}_{it}\mathbf{B} + \mathbf{u}_i + \mathbf{e}_{it} \\ i \in \{1, 2, \ldots, N\},\ \ t \in \{1, 2, \ldots, T_i\} \end{aligned} \tag{1}$$

where $\mathbf{Y}_{it}$ is a $(1 \times k)$ vector of dependent variables, $\mathbf{X}_{it}$ is a $(1 \times l)$ vector of exogenous covariates, and $\mathbf{u_i}$ and $\mathbf{e}_{it}$ are $(1 \times k)$ vectors of dependent variable-specific panel fixed-effects and idiosyncratic errors, respectively. The $(k \times k)$ matrices $\mathbf{A}_1$, $\mathbf{A}_2, \ldots, \mathbf{A}_{p-1}$, $\mathbf{A}_p$ and the $(l \times k)$ matrix $\mathbf{B}$ are parameters to be estimated. We assume that the innovations have the following characteristics: $\mathbf{E}(\mathbf{e}_{it}) = \mathbf{0}$, $\mathbf{E}(\mathbf{e}'_{it}\mathbf{e}_{it}) = \mathbf{\Sigma}$, and $\mathbf{E}(\mathbf{e}'_{it}\mathbf{e}_{is}) = \mathbf{0}$ for all $t > s$.

Similar to Holtz-Eakin, Newey, and Rosen (1988), we assume that the cross-sectional units share the same underlying data generating process, with the reduced-form parameters $\mathbf{A}_1, \mathbf{A}_2, \ldots, \mathbf{A}_{p-1}, \mathbf{A}_p$, and $\mathbf{B}$ to be common among them. Systematic cross-sectional heterogeneity is modeled as panel-specific fixed effects. This setup contrasts with time-series VAR, where by construction, the parameters are specific to the unit being studied, or with random-coefficient panel VAR, where the parameters are estimated as a distribution.[2]

The parameters above may be estimated jointly with the fixed effects or, alternatively, with ordinary least squares (OLS) but with the fixed effects removed after some transformation on the variables. However, with the presence of lagged dependent variables in the right-hand side of the system of equations, estimates would be biased even with large N (Nickell 1981). Although the bias approaches zero as T gets larger, simulations by Judson and Owen (1999) find significant bias even when $T = 30$.

2.1 GMM estimation

Various estimators based on GMM have been proposed to calculate consistent estimates of the above equation, especially in fixed T and large N settings.[3] With our assumption that errors are serially uncorrelated, the model in first difference (FD) may be consistently estimated equation by equation by instrumenting lagged differences with

2. See Canova and Ciccarelli (2013) for a survey of random-coefficient panel VAR models.
3. Other methods include analytical bias correction for the least-squares dummy variable model (for example, Kiviet [1995] and Bun and Carree [2005]) and bias correction based on bootstrap methods (for example, Everaert and Pozzi [2007]).

differences and levels of $\mathbf{Y}_{it}$ from earlier periods as proposed by Anderson and Hsiao (1982). This estimator, however, poses some problems. The FD transformation magnifies the gap in unbalanced panels. For instance, if some $\mathbf{Y}_{it-1}$ are not available, then the FDs at time t and $t-1$ are likewise missing. Also, the necessary time periods each panel is observed gets larger with the lag order of the panel VAR. As an example, for a second-order panel VAR, instruments in levels require that $T_i \geq 5$ realizations be observed for each subject.

Arellano and Bover (1995) proposed forward orthogonal deviation (FOD) as an alternative transformation, which does not share the weaknesses of the FD transformation. Instead of using deviations from past realizations, it subtracts the average of all available future observations, thereby minimizing data loss. Because past realizations are not included in this transformation, they remain valid instruments. Potentially, only the most recent observation is not used in estimation. In a second-order panel VAR, for instance, only $T_i \geq 4$ realizations are necessary to have instruments in levels.

We can improve efficiency by including a longer set of lags as instruments. However, this generally has the unattractive property of reducing observations especially with unbalanced panels or with missing observations. As a remedy, Holtz-Eakin, Newey, and Rosen (1988) proposed creating instruments using available data and substituting missing observations with zero based on the standard assumption that the instruments are uncorrelated with the errors. However, overfitting may be an issue, especially when the time dimension is small, as the GMM estimates approach those from OLS. In such cases, Roodman (2009) advocates reporting the number of instruments used and checking how robust the results are to its reduction. In large N and large T settings, this may not be too much of an issue because the Nickell bias from OLS tends to zero as T tends to infinity (Alvarez and Arellano 2003).

The estimators by Anderson and Hsiao (1982) and by Arellano and Bover (1995), as well as other dynamic panel GMM estimators using similar moment restrictions, like those by Arellano and Bond (1991) and by Blundell and Bond (1998), are designed for "small T, large N" panels, that is, many cross-sectional units observed for few time periods.[4] However, simulations by Judson and Owen (1999) show that the Anderson-Hsiao and the Arellano–Bond estimators perform well even as the time dimension is increased. Alvarez and Arellano (2003) establish for the univariate first-order autoregressive model that the GMM estimators based on orthogonal deviations are N consistent when both N and T tend to infinity but T/N tends to a positive constant that is less than or equal to 2.

In the time-series VAR, it is common to test each variable for stationarity using unit-root tests. This is also relevant in GMM estimation of linear dynamic panel models. As noted by Blundell and Bond (1998) in the univariate case, the GMM estimators suffer from the weak instruments problem when the variable being modeled is near unit root.[5]

4. Roodman (2009) provides an excellent discussion of GMM estimation in a dynamic panel setting and its applications using Stata.
5. When the series has unit root, both FD and FOD transformations leave only the idiosyncratic noise in the data; thus the instruments in levels will be uninformative. See Bond (2002) for a discussion.

The moment conditions become completely irrelevant when unit root is present. As with time-series VAR, pretransforming the variables using growth rates or by differencing may mitigate this problem.

While equation-by-equation GMM estimation yields consistent estimates of panel VAR, fitting the model as a system of equations may result in efficiency gains (Holtz-Eakin, Newey, and Rosen 1988). Suppose the common set of $L \geq kp + l$ instruments is given by the row vector $\mathbf{Z}_{it}$, where $\mathbf{X}_{it} \in \mathbf{Z}_{it}$, and equations are indexed by a number in superscript. Consider the following transformed panel VAR model based on (1) but represented in a more compact form,

$$\begin{aligned}
\mathbf{Y}^*_{it} &= \widetilde{\mathbf{Y}}^*_{it}\mathbf{A} + \mathbf{e}^*_{it} \\
\mathbf{Y}^*_{it} &= \begin{bmatrix}\mathbf{y}^{1*}_{it} & \mathbf{y}^{2*}_{it} & \dots & \mathbf{y}^{k-1*}_{it} & \mathbf{y}^{k*}_{it}\end{bmatrix} \\
\widetilde{\mathbf{Y}^*_{it}} &= \begin{bmatrix}\mathbf{Y}^*_{it-1} & \mathbf{Y}^*_{it-2} & \dots & \mathbf{Y}^*_{it-p+1} & \mathbf{Y}^*_{it-p} & \mathbf{X}^*_{it}\end{bmatrix} \\
\mathbf{e}^*_{it} &= \begin{bmatrix}\mathbf{e}^{1*}_{it} & \mathbf{e}^{2*}_{it} & \dots & \mathbf{e}^{k-1*}_{it} & \mathbf{e}^{k*}_{it}\end{bmatrix} \\
\mathbf{A}' &= \begin{bmatrix}\mathbf{A}'_1 & \mathbf{A}'_2 & \dots & \mathbf{A}'_{p-1} & \mathbf{A}'_p & \mathbf{B}'\end{bmatrix}
\end{aligned}$$

where the asterisk denotes some transformation of the original variable. If we denote the original variable as m_{it}, then the FD transformation implies that $m^*_{it} = m_{it} - m_{it-1}$, while for the forward orthogonal deviation, $m^*_{it} = (m_{it} - \overline{m_{it}})\sqrt{T_{it}/(T_{it}+1)}$, where T_{it} is the number of available future observations for panel i at time t and $\overline{m_{it}}$ is the average of all available future observations.

Suppose we stack observations over panels then over time. The GMM estimator is given by

$$\mathbf{A} = (\widetilde{\mathbf{Y}^*}'\mathbf{Z}\widehat{\mathbf{W}}\mathbf{Z}'\widetilde{\mathbf{Y}^*})^{-1}(\widetilde{\mathbf{Y}^*}'\mathbf{Z}\widehat{\mathbf{W}}\mathbf{Z}'\mathbf{Y}^*) \tag{2}$$

where $\widehat{\mathbf{W}}$ is an $(L \times L)$ weighting matrix assumed to be nonsingular, symmetric, and positive semidefinite. Assuming that $\mathbf{E}(\mathbf{Z}'\mathbf{e}) = \mathbf{0}$ and rank $\mathbf{E}(\widetilde{\mathbf{Y}^*_{it}}'\mathbf{Z}) = kp + l$, the GMM estimator is consistent. The weighting matrix $\widehat{\mathbf{W}}$ may be selected to maximize efficiency (Hansen 1982).

Joint estimation of the system of equations makes cross-equation hypothesis testing straightforward. Wald tests about the parameters may be implemented based on the GMM estimate of $\mathbf{A}$ and its covariance matrix. Granger causality tests, with the hypothesis that all coefficients on the lag of variable m are jointly zero in the equation for variable n, may likewise be carried out using this test.

2.2 Model selection

Panel VAR analysis is predicated upon choosing the optimal lag order in both panel VAR specification and moment condition. Andrews and Lu (2001) proposed MMSC for GMM models based on Hansen's (1982) J statistic of overidentifying restrictions. Their proposed MMSC are analogous to various commonly used maximum likelihood-based model-selection criteria, namely, the Akaike information criteria (AIC) (Akaike 1969), the Bayesian information criteria (BIC) (Schwarz 1978; Rissanen 1978; Akaike 1977), and the Hannan–Quinn information criteria (HQIC) (Hannan and Quinn 1979).

If we apply Andrews and Lu's (2001) MMSC to the GMM estimator in (2), their proposed criteria select the pair of vectors (p, q) that minimizes

$$\text{MMSC}_{\text{BIC},n}(k,p,q) = J_n(k^2p, k^2q) - (|q| - |p|)k^2 \ln n$$

$$\text{MMSC}_{\text{AIC},n}(k,p,q) = J_n(k^2p, k^2q) - 2k^2(|q| - |p|)$$

$$\text{MMSC}_{\text{HQIC},n}(p,q) = J_n(k^2p, k^2q) - Rk^2(|q| - |p|) \ln \ln n \qquad R > 2$$

where $J_n(k,p,q)$ is the J statistic of overidentifying restriction for a k-variate panel VAR of order p and moment conditions based on q lags of the dependent variables with sample size n.

By construction, the above MMSC are available only when $q > p$. As an alternative criterion, the overall coefficient of determination (CD) may be calculated even with just-identified GMM models. Suppose we denote the $(k \times k)$ unconstrained covariance matrix of the dependent variables by Ψ. CD captures the proportion of variation explained by the panel VAR model and may be calculated as

$$\text{CD} = 1 - \frac{\det(\Sigma)}{\det(\Psi)}$$

2.3 Impulse–response

Without loss of generality, we drop the exogenous variables in our notation and focus on the autoregressive structure of the panel VAR in (1). Lütkepohl (2005) and Hamilton (1994) both show that a VAR model is stable if all moduli of the companion matrix $\overline{\mathbf{A}}$ are strictly less than one, where the companion matrix is formed by

$$\overline{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \cdots & \mathbf{A}_p & \mathbf{A}_{p-1} \\ \mathbf{I}_k & \mathbf{O}_k & \cdots & \mathbf{0}_k & \mathbf{0}_k \\ \mathbf{0}_k & \mathbf{I}_k & \cdots & \mathbf{0}_k & \mathbf{0}_k \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_k & \mathbf{0}_k & \cdots & \mathbf{I}_k & \mathbf{0}_k \end{bmatrix}$$

Stability implies that the panel VAR is invertible and has an infinite-order vector moving-average (VMA) representation, providing known interpretation to estimated impulse–response functions (IRFs) and forecast-error variance decompositions (FEVDs). The simple IRF Φ_i may be computed by rewriting the model as an infinite VMA, where Φ_i are the VMA parameters.

$$\mathbf{\Phi}_i = \begin{cases} \mathbf{I}_k & i = 0 \\ \sum_{j=1}^{i} \mathbf{\Phi}_{t-j}\mathbf{A}_j & i = 1, 2, \ldots \end{cases}$$

However, the simple IRFs have no causal interpretation. Because the innovations $\mathbf{e}_{it}$ are correlated contemporaneously, a shock on one variable is likely to be accompanied by shocks in other variables. Suppose we have a matrix $\mathbf{P}$, such that $\mathbf{P}'\mathbf{P} = \mathbf{\Sigma}$. Then $\mathbf{P}$ may be used to orthogonalize the innovations as $\mathbf{e}_{it}\mathbf{P}^{-1}$ and to transform the VMA parameters into the orthogonalized impulse–responses $\mathbf{P}\mathbf{\Phi}_{\boldsymbol{i}}$. The matrix $\mathbf{P}$ effectively imposes identification restrictions on the system of dynamic equations. Sims (1980) proposed the Cholesky decomposition of $\mathbf{\Sigma}$ to impose a recursive structure on a VAR. The decomposition, however, is not unique but depends on the ordering of variables in $\mathbf{\Sigma}$.

IRF confidence intervals may be derived analytically based on the asymptotic distribution of the panel VAR parameters and the cross-equation error variance–covariance matrix. Alternatively, the confidence interval may likewise be estimated using Monte Carlo simulation and bootstrap resampling methods.[6]

2.4 FEVD

The h-step ahead forecast error can be expressed as

$$\mathbf{Y}_{it+h} - \mathbf{E}(\mathbf{Y}_{it+h}) = \sum_{i=0}^{h-1} \mathbf{e}_{i(t+h-i)}\mathbf{\Phi}_i$$

where $\mathbf{Y}_{it+h}$ is the observed vector at time $t+h$ and $\mathbf{E}(\mathbf{Y}_{it_h})$ is the h-step ahead predicted vector made at time t. As with IRFs, we orthogonalize the shocks using the matrix $\mathbf{P}$ to isolate each variable's contribution to the forecast-error variance. The orthogonalized shocks $\mathbf{e}_{it}\mathbf{P}^{-1}$ have a covariance matrix $\mathbf{I}_k$, which allows straightforward decomposition of the forecast-error variance. More specifically, the contribution of a variable m to the h-step ahead forecast-error variance of variable n may be calculated as

$$\sum_{i=0}^{h-1} \boldsymbol{\theta}_{mn}^2 = \sum_{i=1}^{h-1} (\mathbf{i}_n' \mathbf{P}\mathbf{\Phi}_i' \mathbf{i}_m)^2$$

6. See, for instance, Lütkepohl (2005) for details applied in time-series VAR.

where $\mathbf{i}_s$ is the sth column of $\mathbf{I}_k$. In application, the contributions are often normalized relative to the h-step ahead forecast-error variance of variable n,

$$\sum_{i=0}^{h-1} \boldsymbol{\theta}_n^2 = \sum_{i=1}^{h-1} \mathbf{i}_n' \boldsymbol{\Phi}_i' \boldsymbol{\Sigma} \boldsymbol{\Phi}_i \mathbf{i}_n$$

Similar to those of IRFs, confidence intervals may be derived analytically or estimated using various resampling techniques.

3 The commands

Model selection, estimation, and inference about the homogeneous panel VAR model above can be implemented with the new commands **`pvar`**, **`pvarsoc`**, **`pvargranger`**, **`pvarstable`**, **`pvarirf`**, and **`pvarfevd`**. The syntax and outputs are closely patterned after Stata's built-in `var` commands to easily switch between panel and time-series VAR. We describe the commands's syntax in this section and provide examples in section 4.

3.1 pvar

`pvar`[7] fits homogeneous panel VAR models by fitting a multivariate panel regression of each dependent variable on lags of itself, lags of all other dependent variables, and lags of exogenous variables, if any. The estimation is by GMM. The command is implemented using the interactive version of Stata's `gmm` command with analytic derivatives.

Syntax

`pvar` *depvarlist* [*if*] [*in*] [, *options*]

Options

`lags(`#`)` specifies the maximum lag order # to be included in the model. The default is to use the first lag of each variable in *depvarlist*.

`exog(`*varlist*`)` specifies a list of exogenous variables to be included in the panel VAR.

`fod` and `fd` specify how the panel-specific fixed effects will be removed. `fod` specifies that the panel-specific fixed effects be removed using forward orthogonal deviation or Helmert transformation. By default, the first # lags of transformed *depvarlist* in the model are instrumented by the same lags in levels (that is, untransformed). `fod` is the default option. `fd` specifies that the panel-specific fixed effects be removed using first difference instead of forward orthogonal deviations. By default, the first

7. This version of the software corrects the implementation of forward orthogonal deviation used in the earlier version of the program. See, for instance, Head, Lloyd-Ellis, and Sun (2016).

lags of transformed (that is, differenced) *depvarlist* in the model are instrumented by the (#+1)th to (2#)th lags of *depvarlist* in levels (that is, untransformed).

`td` subtracts from each variable in the model its cross-sectional mean before estimation. This could be used to remove common time fixed effects from all the variables prior to any other transformation.

`instlags(`*numlist*`)` overrides the default lag orders of *depvarlist* used as instruments in the model (see the `fod` and `fd` options above that describe which lags are used as default). Instead, *numlist*th lags are used as instruments.

`gmmstyle` specifies that "GMM-style" instruments as proposed by Holtz-Eakin, Newey, and Rosen (1988) be used. Lag length to be used as instruments must be specified with `instlags()`. For each instrument based on lags of *depvarlist*, missing values are substituted with zero. Observations with no valid instruments are excluded. This option is available only with `instlags()`.

`gmmopts(`*options*`)` overrides the default `gmm` options run by `pvar`. Each equation in the model may be accessed individually using the variable names in *depvarlist* as equation names. See [R] **gmm** for the available options.

`vce(`*vcetype* [`, independent`]`)` specifies the type of standard error reported, which includes types that are robust to some types of misspecification, that allow for intragroup correlation, and that use bootstrap or jackknife methods.

vcetype may be `robust`, `cluster` *clustervar*, `bootstrap`, `jackknife`, `hac` *kernel lags*, or `unadjusted`; the default is `vce(unadjusted)`.

`overid` specifies that Hansen's J statistic of overidentifying restriction be reported. This option is available only for overidentified systems.

`level(#)` specifies the confidence level, as a percentage, to be used for reporting confidence intervals. The default is `level(95)` or as set by `set level`.

`noprint` suppresses printing of the coefficient table.

Stored results

`pvar` stores the following in `e()`:

Scalars	
`e(N)`	number of observations
`e(n)`	number of panels
`e(tmin)`	first time period in sample
`e(tmax)`	last time period in sample
`e(tbar)`	average time periods among panels
`e(mlag)`	maximum lag order in panel VAR
`e(N_clust)`	number of clusters
`e(Q)`	criterion function
`e(J)`	Hansen's J chi-squared statistic
`e(J_df)`	J-statistic degrees of freedom
`e(rank)`	rank of `e(V)`
`e(ic)`	number of iterations used by iterative GMM estimator
`e(converged)`	`1` if converged, `0` otherwise
Macros	
`e(cmd)`	`pvar`
`e(cmdline)`	command as typed
`e(depvar)`	names of dependent variables
`e(exog)`	names of exogenous variables, if specified
`e(clustvar)`	name of cluster variable
`e(instr)`	instruments
`e(eqnames)`	equation names
`e(timevar)`	name of time variable
`e(panelvar)`	name of panel variable
`e(properties)`	`b V`
Matrices	
`e(b)`	coefficient vector
`e(V)`	variance–covariance matrix of the estimator
`e(Sigma)`	variance–covariance matrix of the model residuals
`e(W)`	weight matrix used for final round of estimation
`e(init)`	initial values of the estimators
Functions	
`e(sample)`	mark estimation sample

3.2 pvarsoc

`pvarsoc` provides various summary measures to aid the process of panel VAR model selection. It reports the model overall CD, Hansen's (1982) J statistic and corresponding p-value, and MMSC developed by Andrews and Lu (2001) based on the J statistic. Andrew and Lu's criteria are all based on Hansen's J statistic, which requires the number of moment conditions to be greater than the number of endogenous variables in the model. `pvarsoc` uses the estimation sample of the least restrictive panel VAR model, that is, with the highest lag order used, for all models that would be fit by the program.

Syntax

`pvarsoc` *depvarlist* [*if*] [*in*] [, *options*]

Options

`maxlag(#)` specifies the maximum lag order for which the statistics are obtained.

`pinstlag(`*numlist*`)` specifies that *numlist*th lag from the highest lag order of *depvarlist* specified in the panel VAR model implemented using `pvar` be used. This option cannot be specified with the `pvaropts(instlag(`*numlist*`))` option.

`pvaropts(`*options*`)` passes arguments to `pvar`. All arguments specified in *options* are passed to and used by `pvar` in estimation.

Stored results

`pvarsoc` stores the following in `r()`:

Scalars	
`r(N)`	number of observations
`r(n)`	number of panels
`r(tmin)`	first time period in sample
`r(tmax)`	last time period in sample
`r(tbar)`	average time periods among panels
`r(maxlag)`	maximum lag order in panel VAR
Macros	
`r(endog)`	names of endogenous variables
`r(exog)`	names of exogenous variables, if specified
Matrices	
`r(stats)`	CD, J, and p-value, MBIC, MAIC, and MQIC

3.3 pvargranger

The postestimation command `pvargranger` performs Granger causality Wald tests for each equation of the underlying panel VAR model. It provides a convenient alternative to Stata's built-in `test` command.

Syntax

`pvargranger` [, `estimates(`*estname*`)`]

Option

`estimates(`*estname*`)` requests that `pvargranger` use the previously obtained set of panel VAR estimates saved as *estname*. By default, `pvargranger` uses the active (that is, the latest) results.

Stored results

`pvargranger` stores the following in `r()`:

Matrix
`r(pgstats)` chi-squared, degrees of freedom, and p-values

3.4 pvarstable

The postestimation command `pvarstable` checks the stability condition of panel VAR estimates by calculating the modulus of each eigenvalue of the fitted model. Lütkepohl (2005) and Hamilton (1994) both show that a VAR model is stable if all moduli of the companion matrix are strictly less than one. Stability implies that the panel VAR is invertible and has an infinite-order VMA representation, providing known interpretation to estimated IRFs and FEVDs.

Syntax

`pvarstable` [, *options*]

Options

`estimates(`*estname*`)` requests that `pvarstable` use the previously obtained set of `pvar` estimates saved in *estname*. By default, `pvarstable` uses the active estimation results.

`graph` requests `pvarstable` to draw a graph of the eigenvalue of the companion matrix.

`nogrid` suppresses the polar grid circles on the plotted eigenvalues. This option may be specified only with `graph`.

Stored results

`pvarstable` stores the following in `r()`:

Matrices
`r(Re)` real part of the eigenvalues of the companion matrix
`r(Im)` imaginary part of the eigenvalues of the companion matrix
`r(Modulus)` modulus of the eigenvalues of the companion matrix

3.5 pvarirf

The postestimation command `pvarirf` calculates and plots IRFs. Three types of IRF can be estimated: simple IRF (default), orthogonalized IRF based on Cholesky decomposition, and cumulative IRF. `pvarirf` also calculates dynamic multipliers and cumulative dynamic multipliers for exogenous variables. Confidence bands are estimated using Gaussian approximation based on Monte Carlo draws from the fitted panel VAR model.

Syntax

`pvarirf` [, *options*]

Options

`step(#)` specifies the step (forecast) horizon; the default is 10 periods.

`impulse(`*impulsevars*`)` and `response(`*responsevars*`)` specify the impulse and response variables. Usually, one of each is specified, and one graph is drawn. If multiple variables are specified, a separate subgraph is drawn for each impulse–response combination. If `impulse()` and `response()` are not specified, subgraphs are drawn for all combinations of impulse and response variables.

`porder(`*varlist*`)` specifies the Cholesky ordering of the endogenous variables to be used when estimating orthogonalized IRFs as well as the order of the IRF plots. By default, the order in which the variables were originally specified on the `pvar` command is used. This allows a new set of IRFs with a different order to be produced without reestimating the system.

`oirf` requests that orthogonalized IRFs be estimated. The default is simple IRFs.

`dm` estimates dynamic multipliers for exogenous variables instead of IRFs.

`cumulative` computes cumulative IRFs. This option may be combined with `oirf`.

`mc(#)` requests that # Monte Carlo draws be used to estimate the confidence intervals of the IRFs using Gaussian approximation. The default is not to estimate or plot confidence intervals; that is, # = 0.

`table` displays the calculated IRFs as a table. The default is not to tabulate IRFs.

`level(#)` specifies the confidence level, as a percentage, to be used for computing confidence bands. The default is `level(95)` or as set by `set level`. `level()` is available only when `mc(#)` > 1 is specified.

`dots` requests the display of iteration dots. By default, one dot character is displayed for each iteration. A red "x" is displayed if the iteration returns an error.

`save(`*filename*`)` specifies that the calculated IRFs be saved under the name *filename*.

`byoption(`*by_option*`)` affects how the subgraphs are combined, labeled, etc. This option is documented in [G-3] ***by_option***.

`nodraw` suppresses the display of the estimated IRFs.

Stored results

`pvarirf` stores the following in `r()`:

Scalars

`r(iter)`	Monte Carlo iterations
`r(step)`	forecast horizon

Macros

`r(porder)`	Cholesky order of orthogonalized IRF

3.6 pvarfevd

The postestimation command **pvarfevd** computes FEVD based on a Cholesky decomposition of the residual covariance matrix of the underlying panel VAR model. Standard errors and confidence intervals based on Monte Carlo simulation may be optionally computed.

One should exercise caution in interpreting computed FEVD when exogenous variables are included in the underlying panel VAR model. Contributions of exogenous variables, when included in the panel VAR model, to forecast-error variance are disregarded in calculating FEVD.

Syntax

`pvarfevd` [, *options*]

Options

`step(#)` specifies the step (forecast) horizon; the default is 10 periods.

`impulse(`*impulsevars*`)` and `response(`*responsevars*`)` specify the impulse and response variables for which FEVD are to be reported. If `impulse()` or `response()` is not specified, each endogenous variable is used in turn.

`porder(`*varlist*`)` specifies the Cholesky ordering of the endogenous variables to be used when estimating FEVDs. By default, the order in which the variables were originally specified on the underlying `pvar` command is used.

`mc(#)` requests that # Monte Carlo draws be used to estimate the standard errors and the percentile-based 90% confidence intervals of the FEVDs. Computed standard errors and confidence intervals are not displayed but may be saved as a separate file.

`dots` requests the display of iteration dots. By default, one dot character is displayed for each iteration. A red "x" is displayed if the iteration returns an error.

`save(`*filename*`)` specifies that the FEVDs be saved under the name *filename*. In addition, standard errors and percentile-based 90% confidence intervals are saved when `mc(#)` > 1 is specified.

`notable` requests the table be constructed but not displayed.

Stored results

`pvarfevd` stores the following in `r()`:

Scalars

`r(iter)`	Monte Carlo iterations
`r(step)`	forecast horizon

Macros

`r(porder)`	Cholesky order

4 Examples

We illustrate the `pvar` suite of commands by analyzing the relationship between labor supply and wage rate, which has been previously analyzed by Holtz-Eakin, Newey, and Rosen (1988) in their seminal article on panel VAR. Unlike their original implementation, however, we estimate the regression equations simultaneously.

4.1 Panel study of income dynamics

We use `psidextract` data accessible from Stata. We replicate the reduced-form panel VAR presented as table 2 in Holtz-Eakin, Newey, and Rosen (1988) using observations from 528 males over 1976–1982 from the Panel Study of Income and Dynamics (PSID) data available in Stata. In their original analysis, Holtz-Eakin, Newey, and Rosen (1988) used a sample of 898 males observed between 1968 and 1981 using annual hours of work and annual average hourly earnings. Thus our sample composition and the time period are slightly different from the original article. In this illustration, the log-transformed wage rate (`lwage`) and weeks worked (`lwks`) are assumed to be a function of three lags of each of the variables. We also assume that the coefficients on wage rate and weeks worked are common across the sample and that systematic individual heterogeneity is captured by individual fixed effects. The variable `fem` is a binary variable indicating the sex of the respondent.

```
. webuse psidextract
. generate lwks = ln(wks)
. pvar lwks lwage if fem == 0, lags(3)
Panel vector autoregresssion

GMM Estimation

Final GMM Criterion Q(b) =  1.11e-32
Initial weight matrix: Identity
GMM weight matrix:     Robust
                                                  No. of obs      =      1584
                                                  No. of panels   =       528
                                                  Ave. no. of T   =     3.000
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
lwks						
lwks						
L1.	.0477872	.1816701	0.26	0.793	-.3082796	.4038541
L2.	-.1891446	.1002787	-1.89	0.059	-.3856872	.007398
L3.	-.0694588	.0554891	-1.25	0.211	-.1782155	.0392979
lwage						
L1.	-.0069066	.0249964	-0.28	0.782	-.0558987	.0420855
L2.	-.0206062	.0137029	-1.50	0.133	-.0474633	.0062509
L3.	-.0224254	.0141702	-1.58	0.114	-.0501985	.0053476
lwage						
lwks						
L1.	.3516101	.2541961	1.38	0.167	-.146605	.8498253
L2.	.1322435	.123261	1.07	0.283	-.1093435	.3738306
L3.	.0890408	.063914	1.39	0.164	-.0362283	.2143099
lwage						
L1.	.5894378	.0820801	7.18	0.000	.4285638	.7503119
L2.	.1818445	.0480188	3.79	0.000	.0877293	.2759597
L3.	.1337024	.0367614	3.64	0.000	.0616515	.2057533

```
Instruments : l(1/3).(lwks lwage)
```

After fitting the reduced-form panel VAR, we may want to know whether past values of a variable, say, x, are useful in predicting the values of another variable y, conditional on past values of y, that is, whether x "Granger-causes" y (Granger 1969). This is implemented as separate Wald tests with the null hypothesis that the coefficients on all the lags of an endogenous variable are jointly equal to zero; thus the coefficients may be excluded in an equation of the panel VAR model. The `pvargranger` command provides a convenient wrapper to Stata's built-in `test` command to perform the Granger causality tests. The first result below shows the test on whether the coefficients on the three lags of `lwage` appearing on the `lwks` equation are jointly zero. The null hypothesis that `lwage` does not Granger-cause `lwks` is rejected at the 90% confidence level; however, the hypothesis that `lwks` does not Granger-cause `lwage` is not rejected. The second test labeled `ALL` is with respect to the coefficients of all lags of all endogenous variables other than those of the dependent variable being jointly zero. Because we have only two endogenous variables in the panel VAR model, this test is the same as the first test.

```
. pvargranger
  panel VAR-Granger causality Wald test
    Ho: Excluded variable does not Granger-cause Equation variable
    Ha: Excluded variable Granger-causes Equation variable
```

Equation \ Excluded	chi2	df	Prob > chi2
lwks			
lwage	8.924	3	0.030
ALL	8.924	3	0.030
lwage			
lwks	2.452	3	0.484
ALL	2.452	3	0.484

The coefficients on the reduced-form panel VARs cannot be interpreted as causal influences without imposing identifying restrictions on the parameters. If the fitted VAR model is stable, it can be reformulated as an infinite-order VMA, on which assumptions about the error covariance matrix may be imposed. IRFs and FEVDs have known interpretation when the panel VAR model is stable.

After one fits a panel VAR model with **pvar**, the moduli of the companion matrix based on the estimated parameters may be calculated using **pvarstable**. We conclude that the model is stable because all the moduli are smaller than one.

```
. pvarstable
   Eigenvalue stability condition
```

Eigenvalue Real	Eigenvalue Imaginary	Modulus
.9174187	0	.9174187
.1487883	-.4773783	.500028
.1487883	.4773783	.500028
-.1725133	.3267878	.3695282
-.1725133	-.3267878	.3695282
-.2327437	0	.2327437

```
   All the eigenvalues lie inside the unit circle.
   pVAR satisfies stability condition.
```

A graph of the stability test may be produced by adding the **graph** option; see figure 1. We can see that the model is stable because the roots of the companion matrix are all inside the unit circle.

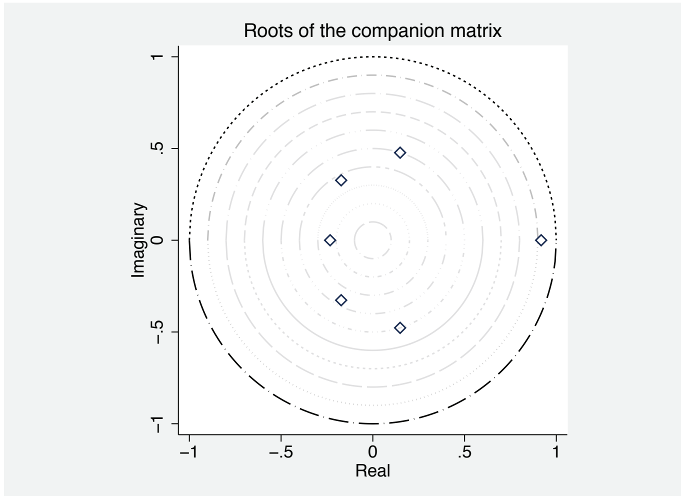

Figure 1. Graph of eigenvalue stability condition

Now that we have established that the panel VAR model is stable, we can calculate IRFs and FEVDs. We can use the `pvarirf` command to calculate simple IRFs, orthogonalized IRFs, cumulative IRFs, and cumulative orthogonalized IRFs. The `pvarfevd` command calculates FEVDs. Orthogonalized IRFs and FEVDs may change depending on how the endogenous variables are ordered in the Cholesky decomposition. Specifically, the ordering constrains the timing of the responses: shocks on variables that come earlier in the ordering will affect subsequent variables contemporaneously, while shocks on variables that come later in the ordering will affect only the previous variables with a lag of one period. Because the ordering of the variable is likely to affect orthogonalized IRFs and the interpretation of the results, one should ensure that the ordering be based on solid theoretical ground. There is no empirical test for the ordering; however, Granger-causality results can be used to add weight to the theoretically chosen ordering. Currently, structural IRFs are not supported, although they may be manually calculated using outputs from `pvar`.

By default, `pvarirf` and `pvarfevd` use the ordering of variables specified in `pvar`. The Cholesky-ordering may be changed easily by using the `porder()` option, instead of reissuing the `pvar` command with the new order of endogenous variables. Confidence intervals are calculated using Monte Carlo simulation.

Following the theoretical exposition by Holtz-Eakin, Newey, and Rosen (1988), we argue that shocks in wage levels have direct impact on contemporaneous hours worked, while current work effort affects wages only in the future. This implies that wages should go first in the order. Note that this ordering is also supported by Granger-causality results reported above: we found that wage Granger-causes weeks, but not vice versa. Using this causal ordering, we calculated the implied orthogonalized IRF using `pvarirf` and the implied FEVD using `pvarfevd`. We used the `porder()` option to put two variables in correct order, which in our example is different from the order listed in `pvar`. Note that ordering does not affect `pvar` estimates; it affects only orthogonalized IRFs and FEVD estimates. The IRF confidence intervals are computed using 200 Monte Carlo draws from the distribution of the fitted reduced-form panel VAR model. Standard errors and confidence intervals for the FEVD estimates are likewise available but not shown here in the interest of space.

```
. pvarirf, oirf mc(200) byoption(yrescale) porder(lwage lwks)
```

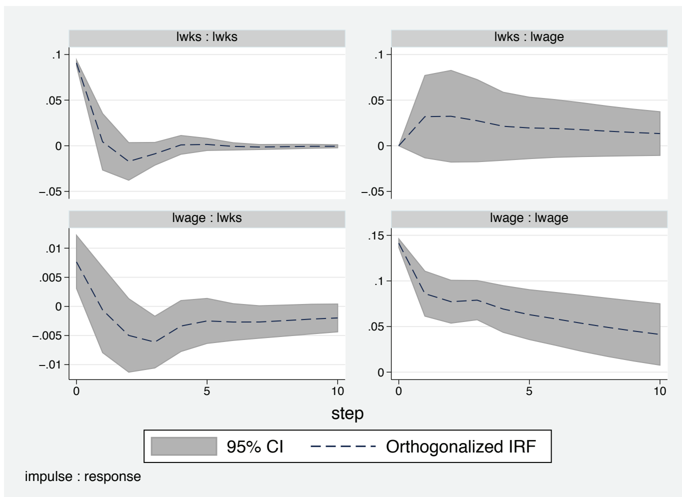

Figure 2. Graphs of orthogonalized IRFs

The IRFs suggest that weeks do not have a significant impact on wage (because the confidence intervals include the zero line in the top right graph of figure 2), while the wage has a nonlinear impact on hours worked: in the first period, it is positive, while in the third period, it turns significantly negative (bottom left graph). Note that the response of wage to weeks is constrained to zero in the first period as a result of the ordering.

```
. pvarfevd, mc(200) porder(lwage lwks) save("fevd_ci.dta")
note: label truncated to 80 characters
Forecast-error variance decomposition
```

Response variable and Forecast horizon	Impulse variable lwage	lwks
lwage		
0	0	0
1	1	0
2	.9641575	.0358426
3	.9417265	.0582735
4	.9335153	.0664847
5	.9312605	.0687395
6	.9296675	.0703325
7	.9279929	.0720071
8	.9266167	.0733833
9	.9256445	.0743555
10	.9249322	.0750679
lwks		
0	0	0
1	.0070398	.9929602
2	.0070685	.9929315
3	.0096815	.9903185
4	.0138484	.9861517
5	.0151277	.9848723
6	.0158209	.9841792
7	.0166287	.9833713
8	.0174266	.9825734
9	.0180863	.9819137
10	.0186139	.9813861

```
FEVD standard errors and confidence intervals based
on 200 Monte Carlo simulations are saved in file
fevd_ci.dta
```

Instead of a priori specifying a third-order panel VAR model, we can use `pvarsoc` to calculate selection-order statistics to identify an optimal moments and model lag order. The command fits `pvar` using preidentified lag orders for the moment instruments and for the panel VAR model and calculates the CD as well as various MMSC if the model is overidentified. It may be necessary to run `pvarsoc` more than once to identify the optimal moment and model lag orders.

Below the order-selection table presents results from the first-, second-, third-, and fourth-order panel VAR models using the first four lags of the endogenous variables as instruments.[8] For the fourth-order panel VAR model, only the CD is calculated because the model is just-identified. Based on the three model-selection criteria by

8. For illustration, we use four lags as instruments to ensure that the GMM model is overidentified for some of the PVAR models. By default, `pvar` runs just-identified models; thus the different MMSC will not be available. We illustrate `instlags()` more extensively in section 4.2 using the National Longitudinal Survey data.

Andrews and Lu (2001), the first-order panel VAR is the preferred model because this has the smallest MBIC, MAIC, and MQIC. While we also want to minimize Hansen's J statistic, it does not correct for the degrees of freedom in the model like the MMSC by Andrews and Lu (2001). Note that the second-order panel VAR models reject Hansen's overidentification restriction at the 5% alpha level, indicating possible misspecification in the model; thus it should not be selected.

```
. pvarsoc lwks lwage if fem == 0, pvaropts(instlags(1/4))
Running panel VAR lag order selection on estimation sample
....
 Selection order criteria
 Sample:  5 - 6                                    No. of obs      =      1056
                                                   No. of panels   =       528
                                                   Ave. no. of T   =     2.000
```

lag	CD	J	J pvalue	MBIC	MAIC	MQIC
1	.9722131	17.13162	.1447131	-66.41531	-6.868385	-29.44043
2	.9830283	18.72182	.0164203	-36.97613	2.721822	-12.32621
3	.9875987	8.959954	.0621083	-18.88902	.9599543	-6.56406
4	.9851995	.	.	.	.	.

We fit the first-order panel VAR model using the first four lags of endogenous variables as instruments because this minimizes each of the MMSC by Andrews and Lu (2001) above. In practice, users should check different sets of lag orders to identify the optimal moment and model lags to be used. We then test for Granger causality and find that we can neither reject the hypothesis that `lwage` does not Granger-cause `lwks` nor reject that `lwks` does not Granger-cause `lwage`.

```
. pvar lwks lwage if fem == 0, lags(1) instlags(1/4)
  (output omitted)
. pvargranger
  panel VAR-Granger causality Wald test
    Ho: Excluded variable does not Granger-cause Equation variable
    Ha: Excluded variable Granger-causes Equation variable
```

Equation \ Excluded	chi2	df	Prob > chi2
lwks			
lwage	0.362	1	0.548
ALL	0.362	1	0.548
lwage			
lwks	0.013	1	0.909
ALL	0.013	1	0.909

The above specifications assume that all the endogenous variables are stationary. The GMM estimator used in `pvar` suffers from weak instrument problems when the variable being modeled is near unit root. The moment conditions become completely irrelevant when the variable has unit root. Using Stata's built-in `xtunitroot` command, we run panel unit-root tests on `lwks` and `lwage` and find that `lwage` has unit root.

```
. xtunitroot ht lwks if fem == 0
. xtunitroot ht lwage if fem == 0
```

We mitigate this issue by using the growth rates of weeks worked, `gwks`, and of wage rate, `gwage`, in the panel VAR model instead of the variables in levels. Another strategy used in the time-series VAR literature when variables have unit roots is to specify the reduced-form VAR model using variables in FDs. Before fitting any model, we test for the presence of unit root in our generated growth-rate variables and find that they are both stationary.

```
. generate gwage = (exp(lwage)-exp(l.lwage))/exp(l.lwage)
(595 missing values generated)
. generate gwks = (wks - l.wks)/l.wks
(595 missing values generated)
. xtunitroot ht gwks if fem == 0
  (output omitted)
. xtunitroot ht gwage if fem == 0
  (output omitted)
. pvarsoc gwks gwage if fem == 0, pvaropts(instlags(1/4))
  (output omitted)
. pvar gwks gwage if fem == 0, lags(1) instlags(1/4)
  (output omitted)
. pvargranger
  panel VAR-Granger causality Wald test
    Ho: Excluded variable does not Granger-cause Equation variable
    Ha: Excluded variable Granger-causes Equation variable
```

Equation \ Excluded	chi2	df	Prob > chi2
gwks			
gwage	0.874	1	0.350
ALL	0.874	1	0.350
gwage			
gwks	0.253	1	0.615
ALL	0.253	1	0.615

Based on the lag-order selection criteria, we fit a first-order panel VAR model using the first four lags of endogenous variables as instruments. We again test for Granger causality and find that in this respecified panel VAR model in growth rates, we can neither reject the hypothesis that `lwage` does not Granger-cause `lwks` nor reject that `lwks` does not Granger-cause `lwage`.

4.2 National Longitudinal Survey

The panel VAR models in the previous section are fit using FODs to remove the individual fixed effects. Another way to remove the fixed effects is to use FDs. Using either FOD or FD should not matter much theoretically when there are no gaps in the data and when there are a large number of cross-sectional units. Gaps in the data, however, are magnified when using FD. Furthermore, in general, FD requires a longer time dimension than FOD, which may be an issue when fitting panel VAR models using short panels.

We illustrate this issue using the subsample of women aged 14–26 years in 1968 from the 1968–1975 National Longitudinal Survey of Youth available from Stata. Holtz-Eakin, Newey, and Rosen (1988) analyzed the 1966–1975 National Longitudinal Survey of Men. As with the earlier examples, we specify homogeneous panel VAR models of log-transformed wage rate (`ln_wage`) and weeks worked (`ln_wks`). Note from the output of `xtdescribe` that in addition to women who were not included in all rounds of the survey, there are two periods when all women have not been observed (shown as dots on the output), representing the years 1974 and 1976.

```
. webuse nlswork2, clear
(National Longitudinal Survey.  Young Women 14-26 years of age in 1968)
. xtdescribe
  idcode:  1, 2, ..., 5159                                  n =       3914
    year:  68, 69, ..., 78                                  T =          9
           Delta(year) = 1 unit
           Span(year)  = 11 periods
           (idcode*year uniquely identifies each observation)
Distribution of T_i:   min      5%     25%       50%       75%     95%     max
                         1       1       2         4         6       9       9
     Freq.  Percent    Cum. |  Pattern
 ---------------------------+-------------
      213      5.44    5.44 |  111111.1.11
      191      4.88   10.32 |  .........11
      173      4.42   14.74 |  .......1.11
      167      4.27   19.01 |  1..........
      134      3.42   22.43 |  ..........1
      116      2.96   25.40 |  ....11.1.11
      113      2.89   28.28 |  .....1.1.11
       93      2.38   30.66 |  ...111.1.11
       93      2.38   33.04 |  ..1111.1.11
     2621     66.96  100.00 | (other patterns)
 ---------------------------+-------------
     3914    100.00         |  XXXXXX.X.XX
. generate ln_wks = ln(wks_work)
(844 missing values generated)
```

Assuming that `ln_wage` and `ln_wks` are stationary, we run first-order panel VAR models using either the `fd` or the `fod` option and using different numbers of lags as instruments. The individual outputs are redacted for conciseness and are instead summarized in one table. Note that in the FD specification, we use the second lags of the untransformed variables as the earliest lag used as instrument, while in the FOD specification, we use the first lag of the untransformed variables. These specifications assume that the original model using untransformed variables have no serial correlation. By construction, first differencing introduces serial correlation in the model; thus only further lags are valid instruments. We present results with two lag options for each model—using one lag (that is, the second in FD and the first in FOD) and using two lags (that is, lags two and three in FD and lags one and two in FOD). The options are specified with the `instl()` option as shown below. When serial correlation is present in the original untransformed model, only more distant lags can be used according to the order of the serial correlation. We run the following commands:

```
. pvar ln_wks ln_wage, fd
. estimates store fd_2
. pvar ln_wks ln_wage, fd instlags(2/3)
. estimates store fd_2t3
. pvar ln_wks ln_wage, fod
. estimates store fod_1
. pvar ln_wks ln_wage, fod instlags(1/2)
. estimates store fod_1t2
```

The table below summarizes the panel VAR models specified above. For now, we focus our attention on the number of observations used in each specification. Note that the FOD specification uses more observations when using either one or two lags as instruments. In both FOD and FD specifications, however, the numbers of observations available for analysis fall when using two lags as instruments, although the drop is bigger for the FD specification because of gaps in the data.

```
. estimates table fd_2 fd_2t3 fod_1 fod_1t2 , b(%3.2f) se(%3.2f)
> stats(N J J_pval) modelwidth(8)
```

Variable	fd_2	fd_2t3	fod_1	fod_1t2
ln_wks				
ln_wks				
L1.	-0.27	0.18	-0.08	0.29
	0.07	0.15	0.05	0.09
ln_wage				
L1.	-0.88	-0.56	-0.32	0.08
	0.16	0.31	0.08	0.12
ln_wage				
ln_wks				
L1.	0.14	0.08	0.15	0.02
	0.03	0.06	0.03	0.04
ln_wage				
L1.	0.50	0.49	0.69	0.65
	0.07	0.12	0.04	0.06
Statistics				
N	3241	1810	4195	2449
J	0.00	29.25	0.00	38.64
J_pval	.	0.00	.	0.00

legend: b/se

The problem of missing observations when using longer lags as instruments may be circumvented by using GMM-style instruments, where missing observations are substituted with zero, as proposed by Holtz-Eakin, Newey, and Rosen (1988). We refit the first-order panel VAR models using either the `fd` or the `fod` option, with two lags of untransformed variables in levels as instruments, but this time specifying the `gmmstyle` to use GMM-style instruments. As a default, observations with missing lagged observations for instruments are dropped. With `gmmstyle` specified, these moment conditions are replaced with zeros and are therefore no longer missing. Here we see that the numbers of observations are the same as when just one lag of the untransformed variable is used as the instrument.

```
. pvar ln_wks ln_wage, fd instlags(2/3) gmmstyle
  (output omitted)
. estimates store fd_2t3g
  (output omitted)
. pvar ln_wks ln_wage, fod instlags(1/2) gmmstyle
  (output omitted)
. estimates store fod_1t2g
  (output omitted)
. estimates table fd_2t3g fod_1t2g, b(%4.2f) se(%4.2f) stats(N J J_pval)
> modelwidth(8)
```

Variable	fd_2t3g	fod_1t2g
ln_wks		
ln_wks		
L1.	-0.17	-0.07
	0.07	0.05
ln_wage		
L1.	-0.74	-0.31
	0.16	0.08
ln_wage		
ln_wks		
L1.	0.12	0.13
	0.03	0.02
ln_wage		
L1.	0.47	0.67
	0.07	0.04
Statistics		
N	3241	4195
J	14.73	14.29
J_pval	0.01	0.01

legend: b/se

In the two sets of estimates above using two lags as instruments, the p-values for the Hansen's J statistics are alarmingly low, indicating some misspecification in the model. One possible issue is that there might be autocorrelation in the model residuals, thereby making the instruments invalid. This may be easily remedied by adjusting the lags used as instruments. For example, using the first three lags of the untransformed variables as instruments in the first- and second-order panel VAR models below gives low p-values for the J statistics.

```
. pvarsoc ln_wks ln_wage, maxl(3) pvaropts(instlags(1/3) fod gmmstyle)
Running panel VAR lag order selection on estimation sample
...
 Selection order criteria
 Sample:  71 - 72                                  No. of obs      =       864
                                                   No. of panels   =       518
                                                   Ave. no. of T   =     1.668

  lag     CD           J        J pvalue     MBIC       MAIC       MQIC

    1   .9924534   23.51504   .0027623  -30.57755   7.515037  -7.065052
    2   .9943494   7.102435    .130573  -19.94386  -.8975646  -8.187609
    3   .9626358          .          .          .          .          .
```

Using the second to the fourth lag of the untransformed variables instead results in more acceptable p-values for the Hansen tests.

```
. pvarsoc ln_wks ln_wage, maxl(3) pvaropts(instlags(2/4) fod gmmstyle)
Running panel VAR lag order selection on estimation sample
...
 Selection order criteria
 Sample:  71 - 72                                  No. of obs      =       864
                                                   No. of panels   =       518
                                                   Ave. no. of T   =     1.668

  lag     CD           J        J pvalue     MBIC       MAIC       MQIC

    1   .9913965   11.11061   .1955104  -42.98197  -4.889392  -19.46948
    2   .9932585   2.925785   .5703209  -24.12051  -5.074215  -12.36426
    3   .7120468          .          .          .          .          .
```

5 Conclusion

In this article, we briefly reviewed panel VAR model selection, estimation, and inference in a GMM framework and introduced a package of commands to fit panel VAR models. We illustrated the commands using two standard Stata datasets.

We conclude with one note of caution to the users of these programs. With the large number of possible combinations of moments and model lags, data transformations, and instrument type that may be implemented, users might be tempted to choose model estimates that fit their expected results. It is always good practice to be up front about the assumptions of the models specified by discussing the set of instruments used, which data transformation is used to remove the fixed effects, etc. And finally, estimates must be checked for robustness to changes in these parameters.

6 Acknowledgments

We acknowledge the valuable comments by Peter Fuleky, Tom Doan, and an anonymous referee.

7 References

Akaike, H. 1969. Fitting autoregressive models for prediction. *Annals of the Institute of Statistical Mathematics* 21: 243–247.

———. 1977. On entropy maximization principle. In *Applications of Statistics*, ed. P. R. Krishnaiah, 27–41. Amsterdam: North-Holland.

Alvarez, J., and M. Arellano. 2003. The time series and cross-section asymptotics of dynamic panel data estimators. *Econometrica* 71: 1121–1159.

Anderson, T. W., and C. Hsiao. 1982. Formulation and estimation of dynamic models using panel data. *Journal of Econometrics* 18: 47–82.

Andrews, D. W. K., and B. Lu. 2001. Consistent model and moment selection procedures for GMM estimation with application to dynamic panel data models. *Journal of Econometrics* 101: 123–164.

Arellano, M., and S. Bond. 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies* 58: 277–297.

Arellano, M., and O. Bover. 1995. Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics* 68: 29–51.

Blundell, R., and S. Bond. 1998. Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics* 87: 115–143.

Bond, S. 2002. Dynamic panel data models: A guide to micro data methods and practice. Working Paper CWP09/02, Cemmap, Institute for Fiscal Studies. http://cemmap.ifs.org.uk/wps/cwp0209.pdf.

Bun, M. J. G., and M. A. Carree. 2005. Bias-corrected estimation in dynamic panel data models. *Journal of Business and Economic Statistics* 23: 200–210.

Canova, F., and M. Ciccarelli. 2013. Panel vector autoregressive Models: A survey. In *Advances in Econometrics: Vol. 32—VAR Models in Macroeconomics—New Developments and Applications: Essays in Honor of Christopher A. Sims*, ed. T. B. Fomby, L. Kilian, and A. Murphy, 205–246. Bingley, UK: Emerald.

Carpenter, S., and S. Demiralp. 2012. Money, reserves, and the transmission of monetary policy: Does the money multiplier exist? *Journal of Macroeconomics* 34: 59–75.

Everaert, G., and L. Pozzi. 2007. Bootstrap-based bias correction for dynamic panels. *Journal of Economic Dynamics and Control* 31: 1160–1184.

Granger, C. W. J. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37: 424–438.

Hamilton, J. D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.

Hannan, E. J., and B. G. Quinn. 1979. The determination of the order of an autoregression. *Journal of the Royal Statistical Society, Series B* 41: 190–195.

Hansen, L. P. 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50: 1029–1054.

Head, A., H. Lloyd-Ellis, and H. Sun. 2014. Search, liquidity, and the dynamics of house prices and construction. *American Economic Review* 104: 1172–1210.

———. 2016. Search, liquidity, and the dynamics of house prices and construction: Corrigendum. *American Economic Review* 106: 1214–1219.

Holtz-Eakin, D., W. Newey, and H. S. Rosen. 1988. Estimating vector autoregressions with panel data. *Econometrica* 56: 1371–1395.

Judson, R. A., and A. L. Owen. 1999. Estimating dynamic panel data models: A guide for macroeconomists. *Economics Letters* 65: 9–15.

Kiviet, J. F. 1995. On bias, inconsistency, and efficiency of various estimators in dynamic panel data models. *Journal of Econometrics* 68: 53–78.

Love, I., and L. Zicchino. 2006. Financial development and dynamic investment behavior: Evidence from panel VAR. *Quarterly Review of Economics and Finance* 46: 190–210.

Lütkepohl, H. 2005. *New Introduction to Multiple Time Series Analysis*. Heidelberg: Springer.

Mora, N., and A. Logan. 2012. Shocks to bank capital: Evidence from UK banks at home and away. *Applied Economics* 44: 1103–1119.

Neumann, T. C., P. V. Fishback, and S. Kantor. 2010. The dynamics of relief spending and the private urban labor market during the new deal. *Journal of Economic History* 70: 195–220.

Nickell, S. 1981. Biases in dynamic models with fixed effects. *Econometrica* 49: 1417–1426.

Rissanen, J. 1978. Modeling by shortest data description. *Automatica* 14: 465–471.

Roodman, D. 2009. How to do xtabond2: An introduction to difference and system GMM in Stata. *Stata Journal* 9: 86–136.

Schwarz, G. 1978. Estimating the dimension of a model. *Annals of Statistics* 6: 461–464.

Sims, C. A. 1980. Macroeconomics and reality. *Econometrica* 48: 1–48.

About the authors

Michael Abrigo is a PhD candidate in the Department of Economics at the University of Hawaii at Manoa and is a research specialist for the Philippine Institute for Development Studies.

Inessa Love is an associate professor in the Department of Economics at the University of Hawaii at Manoa.